D. Relevant publications:

[1] "A Physiological Rule-Based System for Interpreting Pulmonary Function Test Results", J.C. Kunz, R.J. Fallat, D.H. McClung, B.A. Votteri, J.S. Aikins, H.P. Nii, L.M. Fagan, E.A. Feigenbaum, HPP 78-154, Stanford Heuristic Programming Project, 1978.

[2] "Prototypes: An Approach to Knowledge Representation for Hypothesis Formation", Aikins, J.S., HPP-79-10 (working paper), 1979.

[3] "Use of Artificial Intelligence for Interpretation of Physiological Measurements: Pulmonary Function Diagnosis and ICU Ventilator Management", J.C. Kunz, L.M. Fagan, R.J. Fallat, D.H. McClung, Presented at the National Computer Conference by invitation, June 1978.

[4] "Measurement, Modeling and Artificial Intelligence Using Respiratory Data", J.J. Osborn, Keynote Address to Computers in Critical Care and Pulmonary Medicine, May 1979, Proceedings Published by IEEE, 1979.

[5] Automated Interpretation of Respiratory Measurements in the ICU Using Techniques of Artificial Intelligence", J.C. Kunz, L.M. Fagan, E.A. Feigenbaum, J.J. Osborn, Proceedings of Computers in Critical Care and Pulmonary Medicine, IEEE Press, 1979.

[6] "A Physiological Rule-Based System for Interpreting Pulmonary Function Test Results", J.C. Kunz, R.J. Fallat, D.H. McClung, B.A. Votteri, J.S. Aikins, H.P. Nii, L.M. Fagan, E.A. Feigenbaum, Proceedings of Computers in Critical Care and Pulmonary Medicine, IEEE Press, 1979.

[7] "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering", E.A. Feigenbaum, Proceedings of the IJCAI, (1977). (Also Stanford Computer Science Department Memo STAN-CS-77-612).

VM

A. Technical Goals

The Ventilator Manager program (VM) interprets the clinical significance of time varying quantitative physiological data from patients in the ICU. This data is used to manage patients receiving ventilatory assistance. An extension of a physiological monitoring system, VM (1) provides a summary of the patient's physiological status appropriate for the clinician; (2) recognizes untoward events in the patient/machine system and provides suggestions for corrective action; (3) suggests adjustments to ventilatory therapy based on a long-term assessment of the patient status and therapeutic goals; (4) detects possible measurement errors; and, (5) maintains a set of patient-specific expectations and goals for future evaluation. The program produces interpretations of the physiological measurements over time, using a model of the therapeutic procedures in the ICU and clinical knowledge about the diagnostic implications of the data. These therapeutic guidelines are represented by a knowledge base of rules created by clinicians with extensive ICU experience.

The PMC and SUMEX computers will be linked by telephone. The physiological measurements are generated every 2-10 minutes by the PMC computer system. This data now is taken to SUMEX on magnetic tape. It will be provided to VM in real time using the phone link. Information, suggestions to the clinicians, and/or requests for additional information will be sent back to the ICU for action.

B. Medical Relevance and Collaboration

To assist in the interpretation process, VM must be able to recognize unusual or unexpected clinical events (including machine malfunction) in a manner specifically tailored to the patient in question. The interpretation task is viewed as an ongoing process in the ICU, so that the physiological measurements must be continually reevaluated producing a current clinical picture.

This picture can then be compared with previous summary of patient status to recognize changes in patient condition upon which therapy selection and modifications can be made. The program must also determine when the measurements are most likely to be sensitive to error or when external measurements would be of diagnostic significance.

VM offers a new approach towards more accurate recognition of alarm conditions by utilizing the history and situation of the patient in the analysis. This is in contrast to the use of static limits applied to measurements generated to fit the "typical patient" under normal conditions. Our program uses a model of interpretation process, including the types and levels of conclusions drawn manually from the measurements to provide a summary of patient condition and trends. The program generated conclusions are stated at levels more abstract than the raw data; for example, the presence of hemodynamic stability/instability rather than in terms of heart rate and mean arterial pressure. When the data is not reliable enough to make these conclusions, additional tests may be suggested. The recognition of important conclusion for which external verification is sought, will also elicit the suggestion for confirming tests from the program.

C. Progress Summary

VM has been demonstrated using actual patient data recorded on magnetic tape. The input to VM is the values of 30 physiological measurements provided on a 2- or 10-minute bases by a automatic monitoring system. The output is in the form of suggestions to clinicians and periodic summaries (see example case below).

Example Case

The following case demonstrates the current state of development of the system. The data used in this example were obtained from a post-cardiac surgery patient from the ICU at Pacific Medical Center. The terms VOLUME, ASSIST, CONTROLLED MANDATORY VENTILATION (CMV), and T-PIECE refer to specific types of ventilatory assistance. The output format is:(a) ..time of day.., (b) generated comments for clinicians, starting with "**", and (c) commentary in {}.

```
 ..1350.. ..1351..
 ** SYSTEM ASSUMES PATIENT STARTING VOLUME VENTILATION.
                                        {monitoring started}
 ** HYPERVENTILATION                    {diagnostic conclusions
 ** TACHYCARDIA                          based on monitored data}
 ** PATIENT HYPERVENTILATING.           {suggested therapy based on
 ** SUGGEST REDUCING MINUTE VOLUME       diagnosis}
 ..1400..
 . . .
 ..1450..
** HYPERVENTILATION
** TACHYCARDIA
** PATIENT HYPERVENTILATING.
** SUGGEST REDUCING MINUTE VOLUME
 ..1500..
** HYPERVENTILATION
** PATIENT HYPERVENTILATING.
** SUGGEST REDUCING MINUTE VOLUME


Current conclusions:               {summary information}
HYPOTENSION PRESENT for 41 MINUTES
HYPERVENTILATION PRESENT for 33 MINUTES
SYSTOLIC B.P. LOW for 46 MINUTES
{etc.}
Conclusions:       {time of day}   |.....|.....|.....|.
                                   13    14    15    16
HEMODYNAMICS -- STABLE                        ====
HYPERVENTILATION -- PRESENT            =     == == ====
HYPOTENSION -- PRESENT                    ==== ========
TACHYCARDIA -- PRESENT                 ==== ==

patient is on ASSIST                         =====   ==
patient is on CMV                        ====     ===
patient is on VOLUME                   ==
patient is on NOT-MONITORED        ====

Goal is CMV                            ================
Goal is VOLUME                     ====
                                   |.....|.....|.....|.
                                   13    14    15    16
```

The availability of new measurements requires updated interpretations based on the changing values and trends. As the patient setting changes--e.g., as a patient starts to breathe on his own during removal (weaning) from the ventilator--the same measurement values lead to different interpretations. In order to properly interpret data collected during changing therapeutic contexts, the knowledge base includes a model of the stages that a patient follows from admission to the unit through the end of the critical monitoring phase. Recognition of the appropriate patient context is an essential step in determining the meaning of most physiological measurements.

The majority of the knowledge of the VM program is concerned with the relations between the various concepts known by the program. These concepts include: measurement values, typical therapeutic decisions, diagnostic labels, and physiological states. The connections between concepts are represented by a form of production rules using the structure "IF premise THEN action."

The rules in VM are of the form:

IF facts about measurements or previous conclusions are true

THEN

1) Make a conclusion based on these facts;

A sample VM rule is shown below.

```
STATUS RULE: STABLE-HEMODYNAMICS
DEFINITION: Defines stable hemodynamics for most settings
APPLIES to patients on VOLUME, CMV, ASSIST, T-PIECE
COMMENT: Look at mean arterial pressure for changes in
    blood pressure and systolic blood pressure for maximum
    pressures.
IF
  HEART RATE is ACCEPTABLE
  PULSE RATE does NOT CHANGE by 20 beats/minute in 15 minutes
  MEAN ARTERIAL PRESSURE is ACCEPTABLE
  MEAN ARTERIAL PRESSURE does NOT CHANGE by 15 torr in 15 minutes
  SYSTOLIC BLOOD PRESSURE is ACCEPTABLE
THEN
  The HEMODYNAMICS are STABLE
```

Figure 2. Sample VM Interpretation Rule. The meaning of 'ACCEPTABLE' varies with the clinical context--i.e., whether the patient is receiving VOLUME or CMV ventilation, etc. This rule makes a conclusion for internal system use. Similar rules also make suggestions to the user.

II. Research Plans

A. Long Range goals and plans

PUFF

Consensus

Physician acceptance of assistance by knowledge-based programs is understandably inhibited by disagreements between physician diagnoses and those produced by the programs. This disagreement reflects a deeper underlying disagreement among the physicians themselves on the rules to be used for diagnosis. A more subtle problem arises when physicians agree on the diagnosis but cannot agree on the supporting evidence or the reasoning which led to the diagnosis.

We consider obtaining consensus among physicians an important problem for PUFF's acceptance. At the same time, the process of building consensus about a body of knowledge is an interesting area of research in artificial intelligence. The general question to be answered is: What components of the diagnostic process cause differences in the diagnosis? More specifically, in terms of the knowledge which the physicians bring to bear on the diagnostic process, are there differences in representation? (i.e. do physicians use different forms of knowledge?) Are there differences in the definition? (i.e. do physicians use different models or definitions of disease, manifestations, and diagnosis?) Are there differences in the process? (i.e. do physicians reason differently about the problem?)

A second set of questions is: When designing knowledge-based systems, should it be designed to be modifiable to conform to the user physician's diagnostic process? Or, should it be designed to act to identify and record areas of disagreement among the physicians? Should it warn the physicians, mediate the differences, or merely collect the information?

We would like to spend the next few years answering some of these questions in order to understand the problem of consensus building among physicians.

VM

The long range goal of the VM project is to develop and to evaluate an interpretation system which will improve patient care in the ICU.

To insure acceptance by physicians, a careful evaluation of the correctness of the advice of the program against a set of prospective cases will be carried out. A large amount of patient data has already been collected for this evaluation. An experimental period of use in the ICU by clinicians associated with our project is projected during 1979-81.

III. Interactions With The SUMEX-AIM Resource

A. Collaborations and medical use of programs via SUMEX

The PUFF/VM project requires very close collaboration between investigators at two institutions separated by fifty miles. This kind of collaboration, in which program development and testing proceeds concurrently on the same application system, requires a computer network facility for sharing of code, data and ideas. SUMEX has been used at PMC for running programs developed concurrently by Stanford and PMC staff, and data has been taken from the PMC computer system and transferred to SUMEX on magnetic tape for program development and testing.

B. Sharing and interactions with other SUMEX-AIM projects

We have participated in the AIM workshop and had very fruitful interaction with a number of other SUMEX users, directly influencing our perception of important problems and potentially appropriate solutions. Personal contacts at other conferences, at Stanford AI weekly meetings, and at PMC with visiting members of the AIM community, have also been very helpful in keeping abreast of the current thinking of other members of the AI community and with members of the medical community interested in computer based physiological analysis and diagnosis. Specifically, within AIM, there is the closest possible collaboration with researchers on the MYCIN, MOLGEN and DENDRAL projects, who share common space, common techniques, and common attitudes.

C. Critique of resource management

The SUMEX community continues to be an extremely supportive environment in which to do research on uses of artificial intelligence in clinical medicine. The community has two equally vital resources -- the people with knowledge and interest in AI and the facility on which AI system development can proceed. They are equally excellent as resources, helping hands when faced with problems, and friendly support for continued productive research. The availability of INTERLISP; of a facility on which routine data processing functions (eg. manipulating magnetic tapes and making long listings) can take place; and of message-sending among remote users are all vital functions for our project. SUMEX provides them in an environment which is friendly and reliable.

D. Needs and plans for other computational resources

The computation facility at PMC is currently the source of all of the data being used by the PUFF/VM project, and it will continue in this capacity. We expect to link the two machines using a simple telephone dial-up link, but this represents the only system increment to the computational facility of the collaborative project. As the AI techniques developed under PUFF/VM enter routine clinical use at PMC, we have the requirement for system support on which these programs can execute.

VM will enter a period of very rapid development of its knowledge base during the summer of 1979. The basic form of the clinical knowledge has been developed and demonstrated; we expect to work rapidly to fill in a great deal of relevant detail. This knowledge base will require very careful validation as individual subsections near completion. The PUFF/VM grant has plans and resources for bringing large amounts of patient physiological data to SUMEX in real time, and these data will be available starting in the summer of 1979. The process of validation then will require running VM in real time so that PUFF/VM researchers can compare system interpretations of patient state with the actual state as determined by careful concurrent clinical evaluation. We believe that we can effectively use 3-4 hours per day of running VM in a real time test mode during the initial validation period. As the system oration becomes more predictable in 1980, longer running times will be required to identify system problems, and we predict the need to run the system for a full eight hour shift each day.

This test and validation process can be carried out with existing SUMEX resources. The current SUMEX load average is too high to support an effective interaction over a long period of time between clinical researchers and the system. In addition, the very large demand for computation for validating VM will place a great burden on the other SUMEX users. The proposed SUMEX PDP-10-compatible machine will be a potential vehicle for this validation effort if it is available and has easy loading of SUMEX PDP-10 INTERLISP programs onto the new machine.

E. Recommendations for future community and resource development

We perceive the evolution of our AI capability as moving from a highly speculative development state, for which the interactive development capabilities of SUMEX are vital, to a more stable but still changing validation-and-evaluation state. Ultimately we foresee rather stable specification of a program for routine clinical use. Thus, we see the need to transfer our AI techniques from the SUMEX PDP-10 to a local host. For this transfer, a principal long-range need is for software systems that will allow us to run AI systems on a mini-computer after they have been developed on the more powerful SUMEX facility. If the validation of PUFF/VM in the PMC clinical setting shows the programs to be effective in health care, then we hope and expect to be able to provide the capability on a routine basis.

4.1.8 Rutgers Computers in Biomedicine

Rutgers Research Resource - Computers in Biomedicine

Principal Investigator: Saul Amarel
Rutgers University, New Brunswick, New Jersey

I. SUMMARY OF RESEARCH PROGRAM

A) Goals and Approach

The fundamental objective of the Rutgers Resource is to develop a computer based framework for significant research in the biomedical sciences and for the application of research results to the solution of important problems in health care. The focal concept is to introduce advanced methods of computer science - particularly in artificial intelligence - into specific areas of biomedical inquiry. The computer is used as an integral part of the inquiry process, both for the development and organization of knowledge in a domain and for its utilization in problem solving and in processes of experimentation and theory formation.

The Resource community includes 85 researchers and professionals - 37 members, 11 associates, 28 collaborators and 9 users. Members are mainly located at Rutgers. Collaborators are located in several distant sites and they interact, via the SUMEX-AIM and RUTGERS/LCSR facilities, with Resource members on a variety of projects, ranging from system design/improvement to clinical data gathering and testing of expert systems. Our collaborations are described further in section B. below. Resource users are located at Johns Hopkins University, University of Pittsburgh, Stanford University and the NIH campus.

Resource activities include research projects (collaborative research and core research), training/dissemination projects, and computing services in support of pilot user projects. The research projects are organized in three main AREAS OF STUDY. The areas of study and the senior investigators in each of these are:

(1) Medical Modeling and Decision Making (C. Kulikowski).

(2) Modeling Belief Systems and Commonsense Reasoning (C. F. Schmidt and N. S. Sridharan).

(3) Artificial Intelligence: Representations, Reasoning, and System Development (S. Amarel).

The training/dissemination activities of the Rutgers Resource include sponsorship of the Annual AIM Workshop - whose main objective is to strengthen interactions between AIM investigators, to disseminate research methodologies and results, and to stimulate collaborations and imaginative resource sharing within the framework of AIM. The fourth AIM Workshop was held at Rutgers on June 25 to 28, 1978.

The RUTGERS/LCSR computer is being used not only for support of local research projects and AIM Workshop activities; but also for Pilot User Projects in the AIM community - within the general framework of the national AIM project. Computing activities in the Resource are coordinated by S. Levy. The RUTGERS/LCSR facility is directed by C. Hedrick.

B) Medical Relevance; Collaborations:

During 1978-79 we developed a versatile system for building consultation programs, called EXPERT. This system is being used extensively in the development and study of several medical consultation models - in collaboration with clinical investigators from several specialties.

In ophthalmology, the CASNET/Glaucoma knowledge base is being translated into the new EXPERT formalism. The development of the glaucoma knowledge base built in conjunction with the investigators of ONET (ophthalmological network) continues, and is being supplemented by knowledge of Japanese variants of the disease and the decision rules embodying the clinical judgement of a Japanese glaucoma expert: Dr. Y. Kitazawa. Preliminary contacts with British glaucoma specialists have also taken place, raising the prospect that in the near future we will be developing an international knowledge base for this disease. In addition, a model for neuro-ophthalmological consultation is being built in collaboration with Dr. William Hart of the Washington University School of Medicine. In this field we have continued the investigation of anatomical-physiological models for guiding reasoning.

In rheumatology, a vigorous new effort is taking place in developing a rheumatological consultant program using the EXPERT scheme. It has begun with a specialized model of diffuse connective tissue disorders in collaboration with Dr. Gordon Sharp of the University of Missouri at Columbia. Dr. Donald Lindberg, Director of the National Health Care Technology Center at the University of Missouri, initiated this activity and is actively participating and supporting the development of the consultation system. Collaboration with Japanese specialists in this area is also beginning. A knowledge base for generalized consultation in rheumatology to be used by primary care physicians is being developed in conjunction with Dr. William Pincus of the University of California at San Diego, and a similar program is being planned for use in Missouri. These studies are helping us understand better the problems of articulating medical knowledge needed at different levels of detail and specialization and have important implications for the future design of medical knowledge bases. Investigations on the use of a data base of cases for deriving decision-rules, and for updating an existing expert system are also proceeding using rheumatological case records. A rule learning scheme was developed and applied to a data base in allergy and immunology.

Other collaborations have been in the areas of endocrinology, where a thyroid consultation knowledge base was developed in conjunction with Dr. R. A. Nordyke of the Pacific Health Research Institute, and in clinical pathology, where a model for the interpretation of blood chemistry and blood gas analysis data is being developed by Drs. J. Smith and C. Speicher of Ohio State University at Columbus.

All the above applications have shown the versatility of the basic EXPERT representation scheme for rapidly developing medical knowledge bases. By continued testing and development of various domain models, the current boundaries of applicability of the EXPERT formalism are being explored, and new facilities added as required to improve the consultative performance of the programs developed.

In addition to the direct medical collaboration, we have continued investigating problems of modeling in enzyme kinetics with Dr. David Garfinkel of the University of Pennsylvania.

C) Progress Summary:

1. Areas of Study and Projects:

a) Medical Modeling and Decision-Making

Research activities during the past year have concentrated on the development of a new generalized consultative system scheme (knowledge representation and associated strategies of inference), called EXPERT, and its application to a number of different medical domains.

The structure of knowledge in EXPERT involves two data types: findings and hypotheses. The hypotheses (diagnostic, prognostic and treatment selection) are organized as a partially ordered network (PON) using hierarchical and causal relationships. The findings are organized according to observational constraints. Production rules are used to encode inferences among findings, between findings and hypotheses, and among hypotheses. Because of the PON organization of hypotheses, the knowledge base can be pre-compiled with attendant space and time efficiencies in the performance of the consultation programs that call on the knowledge base for decision-making advice.

As described in section B of this report, knowledge bases in ophthalmology, rheumatology, endocrinology and clinical pathology have served to test the versatility of the EXPERT formalism.

Another component of our research involves studies of explicit representations of anatomical and physiological knowledge, and the use of generalized heuristics for reasoning with such models. A preliminary model in neuro-ophthalmology has been developed using the AIMDS and other frame-based representations. In conjunction with the EXPERT investigation this work is yielding insights into the nature and applicability of compiled vs. interpreted medical knowledge. This work has important implications for the practical design of large-scale medical knowledge bases of the future.

Problems of updating a knowledge base and learning decision rules from a data base of case records are two other areas of investigation. A program for rule learning by five different fuzzy-logic heuristic methods was developed and tested using allergy case study data. Problems of the transferability of large-scale consultation programs to a minicomputer environment have also been investigated.

Clinical investigations in thyroid disease and hypertension (by investigators at Pacific Health Research Institute and Johns Hopkins School of Medicine) have been aided by Resource support and development of the BRIGHT system.

b) Modeling of Belief Systems and Commonsense Reasoning

The domain which we have been investigating is the understanding of human actions. A theory of action understanding is fundamental to theories of social psychology, communication and social interaction. Additionally, the understanding of human action involves the use of commonsense concepts of both physical and social causation. Commonsense concepts of causality represent one of the most articulated and complex knowledge bases that we use in understanding our everyday world. In developing a psychological theory of action understanding, we have provided an account of the deep semantic properties that underlie commonsense concepts such as action, plan, belief, and intention. This knowledge is represented in the frame-like AI system, AIMDS, which we developed in the Resource together with a process which exemplifies the use of what we term a <u>hypothesize</u> <u>and</u> <u>revise</u> search strategy. Our psychological theory is embodied in the BELIEVER system.

During this past year our efforts have focused on the investigation of the hypothesize and revise process and, in conjunction with this process, the development of a plan generator that can operate in the context of this process. The hypothesize and revise process is both predictive and responsive to the developing observations of actions and able to function robustly in the "open world" of everyday observations. Such a process is made possible because of the highly articulated nature of the commonsense model of social causality.

As a result of empirically testing the BELIEVER knowledge-based theory we have come to a deeper understanding of the methodology involved in testing such theories. The insights concerning problems of hypothesis formation and test can be seen as a generalization of certain central aspects of the theory of action understanding. This is not surprising since we view certain aspects of hypothesis formation and many aspects of theory testing as a type of interpretation task. As a result of this understanding, we have begun to design a system for the generation and test of hypotheses in our domain of social psychology which formalizes and generalizes upon our experience in testing the BELIEVER theory of action understanding.

c) Artificial Intelligence; Representations, Reasoning and Systems Development

A major part of our effort in this core area continued to be directed to collaborations with investigators in the other applications - oriented projects of the Resource. These collaborations are having an impact on the application areas of the Resource, and they are stimulating work on basic AI issues that are related to designs of knowledge-based systems.

During this period, work on AIMDS has provided a strong focus for studies in knowledge representation and in processes of interpretation and theory formation. This work is continuing in the context of an intensive collaboration towards the development of the BELIEVER theory of action interpretation. As part of this work, considerable progress was made in representations of plan structures and in processes for generating and recognizing plans. Also, the AIMDS system has been extended considerably on basis of experience which has been accumulating from attempts to apply it in new tasks of interpretation, knowledge acquisition/learning and hypothesis formation.

Our efforts to develop unifying principles and conceptual frameworks for the design of AI systems were directed this period to three areas: the analysis of several major interpretation/diagnostic systems developed within the AIM community (MYCIN, CASNET/Glaucoma, BELIEVER and DENDRAL); the study of methods for automatic improvement of problem solving systems via shifts in problem representation obtained by the formation of macromoves from elementary moves; and continuing research on theory formation processes - both in the context of model-guided program formation, and in other rule-learning tasks.

Our research on natural language processing has continued with the objective to develop methods that facilitate communication between people (domain experts, users, designers) and computers. We have taken a fresh look at the problem of developing a convenient man machine interface for a glaucoma consultation system. Building on our previous work in this area, we have added several novel features to the design of our interface processor.

We are continuing to study problems of language acquisition/learning to gain insight into the general problem of knowledge acquisition in expert AI systems. However, we have shifted emphasis this year to an approach which assumes a more active teacher-learner dialogue in the language acquisition process. This led to the identification of rules that govern such a dialogue, and to the design of acquisition processes that embody these rules.

Our commitment to a strong AI programming environment resulted in improvements of the Rutgers/UCI LISP system, as well as in other systems programming developments. These efforts are strengthening the tools for design and experimentation that are available to Resource investigators on the RUTGERS/LCSR computer facility.

2) AIM Dissemination; Training

The Fourth Annual AIM Workshop was held at the Continuing Education Center at Rutgers University on June 25-28, 1978. There were approximately 90 invited participants who attended. This Workshop concentrated on several issues in the construction and validation of consultation programs in medicine, on certain aspects of AI in psychology, on current problems in AI research that are relevant to AIM applications, and on a number of management and funding issues concerning the SUMEX-AIM facility in particular and the AIM community in general.

For the first time in the Workshop series a sizable number of graduate students were invited to the Workshop. The inclusion of the graduate students was an important step in AIM <u>community building</u>, in that the people most actively involved in "systems building" were able to get together and discuss their work, with the invited talks and panels of the Workshop forming the context. Most of the project reports discussed in the evening sessions were presented by graduate students.

The program for the 1978 Workshop consisted of the following four classes of activities:

(i) Invited talks on basic themes. The conference contained three invited talks on basic issues. Dr. C. Kulikowski spoke on AI and medicine, Dr. S. Amarel on current themes in AI research, and Dr. C. Schmidt spoke on the relation of AI and psychology. The complete texts of these talks are being included in the Proceedings that are now in preparation.

(ii) Panel discussions on technical and management issues. The panel discussions were designed to give contrasting views on several issues of current concern. Dr. N. S. Sridharan and Dr. H. Pople chaired a panel on hypothesis formation and revision, a current issue in the construction of knowledge-based systems. Dr. R. Smith chaired a panel on issues in the design and transfer of knowledge-based systems, which spanned a wide range of hardware, software, and representation problems relating to transferring techniques and systems beyond the first stages of research. Dr. E. Feigenbaum chaired a panel on collaboration and AIM policy, which focused on management issues.

(iii) Working groups (organized during the Workshop). There were six working groups, scheduled to run two in parallel for three periods of the Workshop. The working groups were: commentary on AIM systems, choice of representation for knowledge, issues of validation in building AIM systems, methods of plausible reasoning, handling multiple sources of knowledge, and technology options for the SUMEX resources.

(iv) Informal reports on work in progress, mainly by graduate students. A total of sixteen informal reports were given in the evening sessions, two in parallel.

Prof. Herbert Simon of Carnegie-Mellon University gave the keynote address at the Workshop banquet. The banquet was attended by the Workshop participants, plus a number of guests from the Rutgers academic community.

The Proceedings (that are expected to be issued shortly) will give a detailed summary of the panels and working groups, as well as the text of the invited talks.

3) Pilot AIM Projects

The two pilot projects that started in the last period - BRIGHT (an NIH-sponsored clinical data base system) and MAINSAIL - continued this year. In addition, after the new KL-2050 computer was installed at Rutgers, two new pilot projects started: INTERNIST (extensive system development and testing) and CONGEN (mainly in support of demos).

4) Computing Facilities

Our plans to enhance the computing facilities at Rutgers - and to increase their availability to the AIM community - were implemented in the Fall of 1978. At present, the RUTGERS/LCSR facility includes a DEC KL-2050T configuration running under TOPS-20. This facility is accessible via TYMNET. [It was also accessible via ARPANET until April 1979; we expect that our connection to the ARPANET will be restored soon.]

Computing in the Rutgers Resource continues to be distributed between the RUTGERS/LCSR facility and SUMEX-AIM, with the bulk of computing being done at Rutgers.

D) Up-to-Date List of Publications

Amarel, S. and C. Kulikowski (1972) "Medical Decision Making and Computer Modeling," Proc. of 5th International Conference on Systems Science, Honolulu, January 1972.

Amarel, S. (1974) "Inference of Programs from Sample Computations," Proc. of NATO Advanced Study Institute on Computer Oriented Learning Processes, 1974, Bonas, France.

Amarel, S. (1974) "Computer-Based Modeling and Interpretation in Medicine and Psychology: The Rutgers Research Resource," Proc. on Conference on the Computer as a Research Tool in the Life Sciences, June 1974, Aspen, by FASEB; also appears as Computers in Biomedicine TR-29. June 1974, Rutgers University, also in <u>Computers in Life Sciences</u>, W. Siler and D. Lindberg (eds.), Faseb and Plenum, 1975.

Amarel, S. (1976) Abstract of Panel on "AI Applications in Science and Medicine" in 1976 National Computer Conference Program, N.Y., June 7-10, 1976.

Amarel, S. (1976) "Summary of National Computer Conference panel on AI Applications in Science and Medicine" in SIGART Newsletter and SIGBIO Newsletter, Fall 1976.

Amarel, S. (1977) "Summary of Issues and Open Problems in AI Applications," in Report of panel on Applications of Artificial Intelligence, Amarel (ed.), in Proc. IJCAI-77, MIT, Cambridge, Aug. 1977; S. Amarel organized this panel session for IJCAI-77.

Amarel, S., (1978) "Introduction and Overview for AI in Science and Medicine," in 'Session on AI in Science and Medicine, National Computer Conference 1978,' AFIPS Conference Proceedings, Vol. 47, 1978, AFIPS Press.

Bruce, B. (1972) "A Model for Temporal Reference and its Application in a Question Answering Program," in <u>Artificial Intelligence</u>, Vol. 3, Spring 1972.

Bruce, B. (1973) "A Logic for Unknown Outcomes," Notre Dame Journal of Formal Logic; also appears as Computers in Biomedicine, TM-35, Nov. 1973, Rutgers University.

Bruce, B. (1973) "Case Structure Systems," Proc. 3rd International Joint Conference on Artificial Intelligence (IFCAI), August 1973.

Bruce, B. (1975) "Belief Systems and Language Understanding," <u>Current Trends in the Language Sciences</u>, Sedelow and Sedelow (eds.) Houston, in press.

Chokhani, S. and C.A. Kulikowski (1973) "Process Control Model for the Regulation of Intraocular Pressure and Glaucoma," Proc. IEEE Systems, Man and Cybernetics Conf., Boston, November 1973.

Chokhani, S. (1975) "On the Interpretation of Biomathematical Models Within a Class of Decision-Making Procedures," Ph.D. Thesis, Rutgers University; also Computers in Biomedicine TR-43, May 1973.

Fabens, W. (1972) "PEDAGLOT. A Teaching Learning System for Programming Language," Proc. ACM Sigplan Symposium on Pedagogic Languages, January 1972.

Fabens, W. (1975) "PEDAGLOT and Understanding Natural Language Processing." Proc. of the 13th Annual Meeting of the Assoc. of Computational Linguistics, October 30 - November 1, 1975.

Kulikowski, C.A. and S. Weiss (1972) "Strategies for Data Base Utilization in Sequential Pattern Recognition," Proc. IEEE Conf. on Decision and Control, Symp. on Adaptive Processes, December 1972.

Kulikowski, C.A. and S. Weiss (1973) "An Interactive Facility for the Inferential Modeling of Disease," Proc. 7th Annual Princeton Conf. on Information Sciences and Systems, March 1973.

Kulikowski, C.A. (1973) "Theory Formation in Medicine: A Network Structure for Inference," Proc. International Conference on Systems Science, January 1973.

Kulikowski, C.A., S. Weiss and A. Safir (1973) "Glaucoma Diagnosis and Therapy by Computer," Proc. Annual Meeting of the Assoc. for Research in Vision and Ophthalmology, May 1973.

Kulikowski, C.A. (1973) "Medical Decision-Making and the Modeling of Disease," Proc. First International Conf. on Pattern Recognition, October 1973.

Kulikowski, C.A. (1974) "Computer-Based Medical Consultation - A Representation of Treatment Strategies," Proc. Hawaii International Conf. on Systems Science, Jan. 1974.

Kulikowski, C.A. (1974) "A System for Computer-Based Medical Consultation," Proc. National Computer Conference, Chicago, May 1974.

Kulikowski, C.A. and A. Safir (1975) "Computer-Based Systems Vision Care," Proceedings IEEE Intercon, April 1975.

Kulikowski, C.A. and Trigoboff, M. (1975) "A Multiple Hypothesis Selection System for Medical Decision-Making," Proc. 8th Hawaii International Conf. on Systems.

Kulikowski, C.A. and N.S. Sridharan (1975) "Report on the First Annual AIM Workshop on Artificial Intelligence in Medicine." Sigart Newsletter No. 55, December 1975.

Kulikowski, C.A. (1976) "Computer-Based Consultation Systems as a Teaching Tool in Higher Education," 3rd Annual N.J. Conf. on the use of Computers in Higher Education, March 1976.

Kulikowski, C.A., S. Weiss, A. Safir, et al (1976) "Glaucomal Diagnosis and Therapy by Computer: A Collaborative Network Approach" Proc. of ARVO, April 1976.

Kulikowski, C.A., S. Weiss, M. Trigoboff and A. Safir (1976) "Clinical Consultation and the Representation of Disease Processes: Some AI Approaches," AISB Conferences, Edinburgh, July 1976.

Kulikowski, C.A., S. Weiss, A. Safir, (1976) "The CASNET Consultation System for Glaucoma: Development by Collaborative Computer Networking," Biosciences Communication Seminars, Aspen, Colorado.

Kulikowski, C.A., and A. Safir (1976) "Concepts of Computer-Based Modeling in Optics and Refraction," Proceedings Association for Computing Machinery Conference, Houston, Texas, pp. 110-114.

Kulikowski, C.A., S. Weiss and A. Safir (1977) "Computer Networking in the Development of a Glaucoma Consultation System" EDUCOM Fall Conference.

Kulikowski, C.A. (1977) "Representation of Knowledge for Reasoning in Medical Consultation," Panel Presentation, IJCAI, Boston.

Kulikowski, C. A., (1977) "Problems in the Design of Knowledge Bases for Medical Consultation," Proc. First Annual Symposium on Computer Applications in Medical Care, Washington, October 1977, pp. 70-74.

Kulikowski, C. A., (1977) "Model-Based Consultation Systems," Proc. Symposium on Computer-Assisted Decision-Making, Tarrytown, October 1977.

Kulikowski, C. A. and Weiss, S. M., (1978) "The Evaluation of Performance in Empirical and Theoretical Models of Medical Decision-Making," Proc. IEEE Computer Society Workshop in Pattern Recognition and Artificial Intelligence, Princeton, April 1978.

Kulikowski, C. A., (1978) "Artificial Intelligence Approaches to Medical Consultation," Proc. Fourth Illinois Conference on Medical Information Systems, May 1978.

Kulikowski, C. A. (1978) "Strategies of Glaucoma Treatment Planning," National Computer Conference-78, AFIPS Conference Proceedings, Vol. 47, 1978, AFIPS Press, Anaheim.

LeFaivre, R. and A. Walker (1975) "Rutgers Research Resource on Computers in Biomedicine, H," Sigart Newsletter No. 54, October 1975.

LeFaivre, R. (1976) "Procedural Representation in a Fuzzy Problem-Solving System," Proc. Natl. Computer Conf., New York, June 1976.

LeFaivre, R. (1977) "Fuzzy Representation and Approximate Reasoning," submitted to IJCAI-77, MIT.

Mathew, R., C. Kulikowski, and K. Kaplan (1977) "Multilevel Representation for Knowledge Acquisition in Medical Consultation Systems," Proc. MEDINFO Conference, Toronto.

Mauriello, D. (1974) "Simulation of Interaction Between Populations in Freshwater Phytoplankton," Ph.D. Thesis, Rutgers University 1974.

Morgenstern, M. (1978) "Transferring Technology from Research Systems to Users," Proc. The Third Jerusalem Conference on Information Technology, Jerusalem, August 1978.

Safir, A., S. Podos, C. Kulikowski, S. Weiss, et al (1976) "Glaucoma Consultation by Computer," Scientific Exhibit E-22, Ophthalmology Transactions, American Academy of Ophthl. and Otolar., Vol. 81, No. 4.

Sangster, B. C. (1978) "Natural Language Dialogue with Data Base Systems: Designing for the Medical Environment," Proc. of the Third Jerusalem Conference on Information Technology, August 1978.

Schmidt, C. (1972) "A comparison of source unidimensional, multidimensional and set theoretic models for the prediction of judgements of trail implication," Proc. Eastern Psych. Assoc. Meeting, Boston, April 1972.

Schmidt, C.F. and J. D'Addamio (1973) "A Model of the Common Sense Theory of Intension and Personal Causation," Proc. of the 3rd IJCAI, August 1973.

Schmidt, C.F. and A. Sedlak (1973) "An Understanding of Social Episodes," Proc. of Symposium on Social Cognition, American Psych. Assoc. Convention, Montreal, August 1973.

Schmidt, C.F. (1975) "Understanding Human Action," Proc. Theoretical Issues in Natural Language Processing: An Interdisciplinary Workshop in Computational Linguistics, Psychology, Artificial Intelligence, Cambridge, Mass., June 1975. Also appears as Computers in Biomedicine, TM-47, June 1975, Rutgers University.

Schmidt, C. (1975) "Understanding Human Action: Recognizing the Motives," Cognition and Social Behavior, J.S. Carroll and J. Payne (eds.), New York: Lawrence Earlbaum Associates, in press. Also appears as Computers in Biomedicine, TR-45, June 1975, Rutgers University.

Schmidt, C.F., N.S. Sridharan and J.L. Goodson (1976) "Recognizing plans and summarizing actions," Proceedings of the Artificial Intelligence and Simulation of Behavior Conference, University of Edinburgh, Scotland, July 1976.

Schmidt, C. (1976) "Understanding Human Action: Recognizing the Plans and Motives of Other Persons," in Cognition and Social Behavior, J. Carroll and J. Payne (eds.), Potomac, Maryland: Lawrence Earlbaum Associates, 1976.

Schmidt, C.F. and J.L. Goodson (1976) "The Subjective Organization of Summaries of Action Sequences," 17th Annual Meeting of the Psychonomic Society, St. Louis, 1976.

Schmidt, C.F., N.S. Sridharan, and J. L. Goodson (1977) "Plan Recognition: A Hypothesize and Revise Paradigm" Proceedings of the 5th International Joint Conference on Artificial Intelligence, Cambridge, Mass.

Schmidt, C.F., and N.S. Sridharan (1977), "Plan Recognition Using a Hypothesize and Revise Paradigm: An Example," in Fifth International Joint Conference on Artificial Intelligence: Advance Papers of the Conference, Vol. 1, Cambridge, Mass., MIT, PP.480-486.

Schmidt, C.F, N.S. Sridharan, and J.L. Goodson (1978) "The Plan Recognition Problem: An Intersection Between Psychology and Artificial Intelligence," Artificial Intelligence, Special Issue on Applications to the Sciences and Medicine, Vol. 11, pp. 45-83.

Sridharan, N.S. (1976) "The Frame and Focus Problems in AI: Decision in Relation to the BELIEVER System. Proceedings of the Conference on Artificial Intelligence and the Simulation of Human Behavior, Edinburgh, July 1976.

Sridharan, N.S. (1976) "An Artificial Intelligence System to Model and Guide Organic Chemical Synthesis Planning", in Chemical Synthesis by Computer, American Chemical Society Press, September 1976.

Sridharan, N.S. and F. Hawrusik (1977) "Representation of Actions with Side-Effects," in Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Boston.

Sridharan, N.S. and Smith, D., (1978) "Design for a Plan Hypothesizer," Proceedings of the AISB/GI Conference, Hamburg, July 1978; (also, CBM-TR-85, DCS, Rutgers University).

Sridharan, N.S. and Schmidt, C.F., (1978) "Knowledge-Directed Inference in BELIEVER," In D. Waterman & F. Hayes-Roth (eds.), Pattern Directed Inference Systems, New York, Academic Press, 1978, pp. 360-379.

Sridharan, N.S., (1978) "Guest Editorial," Special Issue on Applications of AI to the Sciences and Medicine, Artificial Intelligence, Vol. 11, No. 1, 2, August 1978.

Srinivasan, C.V. (1973) "The Architecture of a Coherent Information System: A General Problem Solving System," Proc. of the 3rd IJCAI, August 1973.

Trigoboff, M. (1976) Propagation of Information in a Semantic Net," Proc. of the Conference on Artificial Intelligence and the Simulation of Behaviour, Edinburgh, Scotland, July 1976; updated version appears in CBM-TM-57, Dept. of Computer Science, Rutgers University, 1977.

Trigoboff, M. and C. Kulikowski (1977) "IRIS: A System for Propagation of Inference in a Semantic Network," Proc. Fifth International Joint Conference on Artificial Intelligence, Boston.

Trigoboff, M.L. (1978) "IRIS: A Framework for the Construction of Clinical Consultation Systems," Ph.D. Dissertation, Rutgers University, May 1978; (also, CBM-TR-90, DCS, Rutgers University).

Tucker, S.S. (1974) Cobalt Kinetics in Aquatic Microcosms," Ph.D. thesis, Rutgers University.

Van der Mude, A. and A. Walker (1976) "Some Results on the Inference of Stochastic Grammars," abstract in Proc. Symposium on New Directions and Recent Results in Algorithms and Complexity. Dept. of Computer Science, Carnegie-Mellon University; Information and Control, (in press).

Van der Mude, A. and A. Walker (1978) "On the Inference of Stochastic Regular Grammars," Information and Control, (to appear).

Vichnevetsky, R. (1973) "Physical Criteria in the Evaluation of Computer Methods for Partial Differential Equations," Proc. 7th International AICA Congress, Prague, Sept. 1973; reprinted in Proc. of AICA, Vol. XVI, No. 1, Jan. 1974, European Academic Press, Brussels, Belgium.

Vichnevetsky, R., K.W. Tu, and J.A. Steen (1974), "Quantitative Error Analysis of Numerical Methods for Partial Differential Equations," Proc. 8th Annual Princeton Conference on Information Science and Systems, Princeton University, March 1974.

Walker, A. (1975) "Formal Grammars and the Regeneration Capability of Biological Systems," Journal Comp. and Syst. Sciences, Vol. 11, No. 2, 252-261.

Weiss, S. (1974) "A System for Model-Based Computer-Aided Diagnosis and Therapy," Parts I and II, Ph.D. Thesis, Rutgers University; also Computers in Biomedicine TR-27, Feb. 1974.

Weiss, S. (1976) "A System for Interactive Analysis of a Time-Sequenced Ophthalmological Data Base, Proceedings of Third Illinois Conference on Medical Information Systems.

Weiss, S.M., Kulikowski, C.A. and Safir, A., (1978) "Glaucoma Consultation by Computer," Computers in Biology and Medicine, 8: 25-40, January 1978.

Weiss, S.M., Kulikowski, C.A., Amarel, S., and Safir, A., (1978), "A Model-Based Method for Computer-Aided Medical Decision-Making," Artificial Intelligence, Special Issue on Applications to the Sciences and Medicine, (11, 1978); (also, CBM-TR-63, DCS, Rutgers University, April 1978.)

Weiss, S.M., Kulikowski, C.A., Mizoguchi, F. and Kitazawa, V., (1978), "A Computer-Based Comparison of Japanese and American Decision-Making," Proc. International Conference on Cybernetics and Society, Tokyo, November 1978.

Weiss, S.M., Kern, K.B., Kulikowski, C.A. and Pincus, W., (1978), "An Interactive System for the Design of Classifiers in Diagnostic Applications," Proc. International Conference on Cybernetics and Society, Tokyo, November 1978.

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

The main use of the SUMEX-AIM facility by Rutgers Resource investigators is for system testing with collaborators around the country and for communications on collaborative projects - such as the AIM Workshop - with other scientists in the AIM community. Also, the level of interactions among systems people at the Stanford and Rutgers facilities has increased since the RUTGERS/LCSR system has become a 'server' node in the national AIM network. We continue to access SUMEX-AIM via TYMNET, and (until April-78) via ARPANET.

The ONET community works with the CASNET and EXPERT systems at both the Stanford and Rutgers sites; and up-to-date system copies are maintained at both sites. The research with the U. of Missouri on Rheumatic diseases (using EXPERT) is being done at both sites and each is being used to back up the other. Development of the INTERNIST system by the U. of Pittsburgh researchers has been going on on both the Stanford and Rutgers systems; each system backs up the other for use and demos. These network activities have been temporarily disrupted by the recent removal of the ARPANET link to the Rutgers system; we hope, however, to be reconnected to ARPANET in the near future, and these activities are expected then to go back to their previous levels.

SUMEX-AIM continues to provide the focus for our collaboration with Stanford on the AI Handbook project and for many of the activities surrounding the organization and running of the AIM Workshop. Stanford developments, such as the AGE system, are being explored by Rutgers investigators on the SUMEX-AIM facility. In general, SUMEX-AIM continues to play a vital role as a center for the exchange of ideas and systems with other scientists in the AIM community.

III. RESEARCH PLANS (8/79 - 7/81)

A) Long Range Project Goals and Plans

We are planning to continue along the main lines of research that we have established in the Resource to date with emphasis on broadening our activities in the medical systems area. We also plan to continue our participation in AIM dissemination and training activities as well as our contribution - via the RUTGERS/LCSR computer - to the shared computing facilities of the national AIM network.

B) Justification and Requirements for continued SUMEX use

The use of SUMEX-AIM by the Rutgers Resource will continue to be needed in support of our collaborative research projects, as well as joint AIM dissemination projects and for communications. We estimate a modest total annual level of usage at SUMEX-AIM of about 750 connect hours at an average compute to connect ratio of 1:25. Also, the participation of the RUTGERS/LCSR facility as a node in the AIM network of shared resources will continue to require close coupling and coordination with SUMEX-AIM at the systems level.

C) Plans for Computational Resources

The usage of the RUTGERS/LCSR computer is approaching the capacity limits of the current configuration. There is, however, considerable room to extend the system capacity by augmenting key components such as main memory. Also, considerable effort is still required to improve the system's environment and thus its reliability. These are directions of development that we foresee for the computational resources that are located at Rutgers. As far as more general planning for new resources is concerned, we believe that the emphasis should be on relatively small machines (DEC's 2020, VAX or equivalent) to be located in major user sites and to be configured as parts of a national network of AIM shared resources.

4.1.9 Simulation of Cognitive Processes (SCP)

Simulation of Cognitive Processes (SCP)

James G. Greeno and Alan M. Lesgold
Learning Research and Development Center
University of Pittsburgh
Pittsburgh, Pennsylvania

I. SUMMARY OF RESEARCH PROGRAM

Technical Goals

The goals of this project remain the simulation of children's behavior in mathematics and reading tasks, done in such a way that various levels of proficiency, from moderate dysfunction and normal prereading skills to expertise, can be modeled. As this work continues, the investigators and several colleagues have also become involved in a more general study of the course of cognitive skill acquisition. The long-term (one to two years from now) effect of this additional work will be to force a more general level of theorizing about cognitive skills and about the course of learning. Such a general approach must, on the one hand, seek out the basic principles of cognitive learning and of how different levels of prior skill and aptitude interact with different types of training to produce learning. However, the work must be grounded in specific, detailed work on the learning and performance of specific skills.

Our work on SUMEX-AIM continues, therefore, to focus specifically on mathematics and reading. However, we have also begun a new activity, the development of a procedure and language acquisition system (PLAS). PLAS is part of an on-going project to investigate the utility of representing the skills of mathematics children demonstrate as well as the errors they make while learning mathematics in terms of a model of the acquisition of language and procedural knowledge. The system is a version of Anderson's Language Acquisition System, with a modified memory structure that incorporates structured programs as well as declarative knowledge. The system can understand sentences of arithmetic (construct memory structures which are executable as programs), speak (given a memory structure, a sentence in the language can be produced), and learn the meaning and syntax of new sentences (given a sentence-memory pair, the system constructs an augmented transition network to translate between them). Such a system, when given incorrect or incomplete knowledge, will produce behavior inappropriate for the task at hand. It is the purpose of our group to relate specific mathematical errors that people (children and adults) make to simulatable misconceptions and thereby to deduce techniques which will "debug" and enhance such knowledge.

In reading, as our work progresses on a longitudinal study of the course of acquisition for the subprocesses of reading, the task of simulating the interactions of various levels of the reading process is becoming much better defined. It is now clear that if these processes are to be simulated within ACT, as planned originally, it will be necessary for ACT to grow to include partial matching schemes for the testing of production conditions. Thus, we are testing out pieces of a simulation while awaiting the development of the next generation of ACT. Our goals remain the characterization within ACT of the interactions of recognition and comprehension processes in children with different levels of automation of recognition processes, different levels of knowledge of the subject matter domain, and different basic cognitive aptitudes.

Medical Relevance and Collaboration

The range of ability levels we are dealing with in arithmetic and reading includes children who are below average and, in the case of reading, some children who are classified as learning disabled. By providing a framework within which the effects of differing levels of skill acquisition can be understood, we hope to eliminate the spurious use of vague medical categories such as minimal brain damage, etc., and thus more clearly delimit the cases in which there is a real medical problem from other cases in which a person is poor in basic cognitive skills.

Progress Summary

Progress this past year has been greatest in the area of arithmetic, though there has been some progress on reading, also. In arithmetic, work has progressed on children's ability to solve arithmetic word problems and on the PLAS system mentioned above. There has also been some work on geometry. Finally, there has been considerable empirical progress and clarification of our AI goals for the reading work. These will all be discussed in turn.

Arithmetic work

Three reports of work relevant to arithmetic learning have been written (Riley, 1979; Heller, 1979; Greeno, in press). Riley's and Heller's papers present further work investigating the processes that Heller modeled in the ACT system on SUMEX last year. Greeno's paper summarizes issues and results of that work, as well as Danforth's project on learning arithmetic syntax.

Procedure and Language Acquisition System - PLAS

Current status of PLAS. The system being developed is nearing the completion of its component modules. The memory building routines, the derivation tree to ATN grammar translator, and the ATN parser are all complete. There remains the derivation tree constructor (equivalent to BRACKET in LAS but simpler), the building of the module to merge ATNs, and the high level control structure of the system. This work has been aided by informal discussions with John R. Anderson and Patrick W. Langley at Carnegie-Mellon University and the ACT project. There follows a short summary of the most significant parts PLAS: memory, grammar, and syntax acquisition.

The memory of PLAS. All declarative and procedural knowledge of PLAS is stored in a uniform memory structure. There are nodes in the memory which are implemented in UCI LISP as atoms. Each node has a list of other nodes associated with it (has the property AS with a list as its value). The node is called the ASSOCIATION and its list of associated nodes its OFFSPRING (or ASSOCIATIVE). Each node also has a list of dominating nodes or PARENTS (also called the nodes AFFINITIVE) which provide a doubly linked structure for memory searching. There is nothing more to the memory than that. Meaning is provided by linking some of the nodes (terminal nodes) to either executable primitive procedures or words of the language. Complex programs are constructed simply by creating associations of lower level programs. Declarative knowledge is incorporated through consistency of association of words with specific nodes in memory. In this way it is possible to combine declarative and procedural knowledge (not all such associations are useful or executable).

The augmented transition network grammar. The system contains an Augmented Transition Network (ATN) Grammar to parse input sentences and construct memory associations. The grammar is implemented as a finite labeled graph. The vertices of the graph are ATNs and the labels of the edges are actions to be performed (memory structures built) upon successful testing of a precondition. Each precondition is itself an ATN so that the graph representing an ATN is nothing more than a vertex (the graph) which has a set of triples (ATN ACTION ATN) representing the edges leading out of this vertex. Terminal ATNs test for the presence of a given word at the current position of the input string. They either succeed or fail.

Since it is not possible to determine whether a sentence has been successfully parsed until the last word has been examined, the actions on the links are pushed onto an action stack and executed sequentially when the parse succeeds. The actions are nothing more than the creation of associations. In fact the labels of the graph edges are associations in the memory. If a node of the memory is unbound then a new node is created and bound to it. The association thus formed is then available for the next action which may use it or some previously built association. In this way, the distinction bound vs. unbound controls the context of the discourse allowing associations from previous sentences to be incorporated into the meaning of the present sentence. Once the action stack has been processed, the last association is executed, in the sense that associations that contain primitive procedures are executed whereas declarative associations are skipped. Note that it is possible to "talk about" a procedure as well as "do" a procedure if the procedure is embedded in a declarative association.

The learning of syntax. The ATN grammar forces a hierarchical structuring of the parse of any sentence. When a sentence cannot be parsed by the grammar and the sentence is correct, then something else must be done to provide a derivation tree for the sentence. In PLAS the meaning of a new sentence is introduced directly into the system's memory. The only distinction between the associations for this sentence and those of any other sentence previously introduced is that of recency. This is expressed by the order of nodes in an association's AFFINITIVE (parents) list. Given the sentence and the updated memory, an exhaustive search is performed emanating from a specified node to each of the words in the sentence subject to the constraint that the linked structure so formed is a tree. The specified node can be any node of memory. To correspond with LAS, the node is the principal proposition of the meaning just introduced. Note that the node could also be the principal association of the last sentence parsed.

Given the derivation tree of a sentence, PLAS constructs an ATN grammar which will parse the sentence. This "mini-grammar" is then merged with the ATN grammar constructed from previous sentence-memory pairs. The merging process entails the recognition of identical subgrammars and pattern matching between the memory associations (the actions) of the grammar.

Geometry

Two technical reports have been written (Greeno, Magone, and Chaiklin, in press; Greeno, in press) about the geometry work. They report an extension of an earlier model of geometry problem solving developed earlier in ACT (though not on SUMEX). The new development is a simulation of hierarchically organized planning knowledge of the kind analyzed by Sacerdoti. This way of organizing knowledge provides a natural interpretation of the occurrence of constructions. Plan schemata have the form of global actions, and knowledge includes the preconditions needed for each schema to apply. If none of the required preconditions is present in complete form, the system can identify partial matches to some patterns that would provide the needed preconditions and invoke procedures for pattern completion that involve constructing auxiliary lines. The structure also provides a natural interpretation of classical problem-solving set, since the top-down nature of planning can result in adoption of a plan that is then carried out without noticing the possibility of an alternative plan that might be easier.

Reading

The reading work has been mainly empirical this year. We have discovered that reasonable simulation of process interactions in reading will require extensions of ACT to allow partial matching schemes for production condition testing. We have had discussions with Anderson on this and have agreed to continue work on pieces of the simulation within ACTF, while awaiting ACTG, which will permit more natural simulation. Existing fragmentary simulations have allowed clarification of our models of reading dysfunction and permitted sensible redesign of later portions of a longitudinal study of reading acquisition (Lesgold, 1978; Lesgold, Resnick, & Beck, 1979). The next step is partial simulations of local-coherence resolving processes (Lesgold, Roth, & Curtis, in press).

List of Relevant Publications

Greeno, J.G. Constructions in geometry problem solving. LRDC Technical Report, in press.

Greeno, J.G. Preliminary steps toward a cognitive model of learning primary mathematics. In K. Fuson and W. Geeslin (Eds.), Models of children's mathematical learning, ERIC Information Center, in press.

Greeno, J.G., Magone, M.E., & Chaiklin, S. Theory of constructions and set in problem solving. LRDC Technical Report, in press.

Heller, J.I. Schemata in the solution of arithmetic word problems. Paper presented at the annual meeting of the American Educational Research Association, San Francisco, April, 1979.

Lesgold, A.M. Word-memory access and reading: Toward more explanatory evidence. Paper presented at the annual meeting of the American Educational Research Association, Toronto, March, 1978.